

What is Claimed is:

1. A method for distributed lock management comprising:
determining, by a requesting node requesting a lock, a current lock owner of the
5 lock in accordance with an indicator specifying which one or more nodes are available as
lock owners; and
sending a lock request message to said current lock owner; and
wherein
said requesting node waits for an indication as to whether the lock request
10 message has been granted or denied, and said determining is performed by said
requesting node prior to said sending for each request.
2. The method of Claim 1, wherein said lock request message includes a time
period indicated an amount of time for which said requesting node is requesting said
15 lock.
3. The method of Claim 2, further comprising performing, by said current lock
owner:
receiving said lock request message;
20 determining whether said lock is currently granted to another node; and
granting said lock to said requesting node if said lock is not currently granted to
another node, and denying said lock to said requesting node otherwise.

4. The method of Claim 3, further comprising performing, by said current lock owner:

adding said time period to a current time producing an expiration time, said
5 expiration time being used in connection with a subsequent determination as to whether the lock is currently granted to a node.

5. The method of Claim 4, further comprising performing, by said current lock owner:

10 comparing said current time to said expiration time to determine whether said lock is currently granted.

6. The method of Claim 5, further comprising:

sending, by said current lock owner to said requesting node, a lock response
15 message indicating whether said lock request message has been granted or denied.

7. The method of Claim 6, further comprising:

determining, by said requesting node, that said lock request message has been
denied if one of: no lock response message is received from said lock owner in
20 accordance with a timeout condition, and a lock response message is received indicating that said lock request message has been denied.

8. The method of Claim 7, wherein said requesting node is granted said lock until a lock expiration time is reached, said requesting node becomes unavailable, and said lock expires automatically when said lock expiration time is reached.

5 9. The method of Claim 7, wherein said current lock owner becomes unavailable as a lock owner, and the method further comprises performing, by a second node:

 determining a new current lock owner of the lock in accordance with an indicator specifying which one or more nodes are available as lock owners, said indicator indicating that said current lock owner is not available; and

10 sending a lock request message to said new current lock owner, wherein said determining provides for dynamic reassignment of a new lock owner of said lock.

10. The method of Claim 9, wherein said method is performed in a data storage system including a plurality of nodes, each of said plurality of nodes being a director in
15 said data storage system.

11. The method of Claim 10, wherein said lock is associated with a resource used by two or more nodes in said data storage system.

20 12. The method of Claim 11, wherein said resource is a portion of global memory.

13. The method of Claim 11, wherein each of said lock request message and said lock response message are sent using a message switch included in said data storage system.

5 14. The method of Claim 13, wherein said lock request message and said lock response message are transmitted between nodes in said data storage system using portions of a message switch used in transmitting data messages.

10 15. The method of Claim 10, wherein each of said plurality of nodes indicated as being available as a lock owner in accordance with said indicator is designated as a lock owner of a lock.

15 16. The method of Claim 15, wherein any one of said plurality of nodes may request any lock, and wherein any one of said plurality of nodes may be designated as a lock owner of any lock in accordance with said indicator.

20 17. The method of Claim 1, wherein said method is performed in a data storage system including a plurality of nodes, said plurality of nodes being directors that are endpoint in said data storage system, and said determining and said sending are performed by each of said plurality of nodes prior to sending a lock request message.

18. A method for distributed lock management in a data storage system
comprising:

determining, by a requesting node requesting a lock, a current lock owner of the
lock in accordance with an indicator specifying which one or more nodes are available as
5 lock owners; and

sending a lock request message to said current lock owner, and wherein said
determining is performed by said requesting node prior to said sending for each request.

19. The method of Claim 18, wherein said data storage system includes a
10 plurality of nodes which are processors, each of said plurality of nodes being an endpoint
and being a lock owner as designated by said indicator when said each node is available
as a lock holder.

20. The method of Claim 18, wherein said lock request message includes a time
15 period for which said requesting node is requesting said lock, and wherein said current
lock owner determines an automatic expiration period for said lock using said time
period.

21. The method of Claim 20, wherein said current lock owner determines
20 whether to grant or deny said lock request message in accordance with a current value for
said automatic expiration time period determined in accordance with a prior lock request
message.

22. A data storage system comprising:

a plurality of processors which are endpoints in said data storage system;

a message switch used by said plurality of processors to send transmissions, said transmissions including a lock request message and a lock response message;

5 each of said plurality of processors including machine executable instructions on a computer readable medium for processing a lock request from said each processor to another processor that:

determines, prior to sending a lock request message, a current lock owner in accordance with an indicator specifying which one or more of said plurality of

10 processors are available as lock owners; and

sending said lock request message to said current lock owner; and

each of said plurality of processors including machine executable instructions on a computer readable medium for processing a received lock request message when said each processor is indicated as a lock owner by said indicator of another processor that

15 sent said received lock request message that:

determines whether a requested lock in accordance with said received lock request message is currently granted;

granting to said requested lock to a requesting processor if said requested lock is not currently granted, and otherwise denying said requested lock to said requested
20 node, wherein said lock owner determines an automatic expiration time in accordance with a requested lock time included in said received lock request message.

23. A computer program product for distributed lock management comprising:
executable code that determines, by a requesting node requesting a lock, a current
lock owner of the lock in accordance with an indicator specifying which one or more
nodes are available as lock owners; and

5 executable code that sends a lock request message to said current lock owner; and
wherein

said requesting node waits for an indication as to whether the lock request
message has been granted or denied, and said executable code that determines is executed
by said requesting node prior to executing said executable code that sends for each
10 request.

24. The computer program product of Claim 23, wherein said lock request
message includes a time period indicated an amount of time for which said requesting
node is requesting said lock.

15 25. The computer program product of Claim 24, further comprising executable
code that performs, by said current lock owner:

receiving said lock request message;

determining whether said lock is currently granted to another node; and

20 granting said lock to said requesting node if said lock is not currently granted to
another node, and denying said lock to said requesting node otherwise.

26. The computer program product of Claim 25, further comprising executable code that performs, by said current lock owner:

adding said time period to a current time producing an expiration time, said expiration time being used in connection with a subsequent determination as to whether

5 the lock is currently granted to a node.

27. The computer program product of Claim 26, further comprising executable code that performs by said current lock owner:

comparing said current time to said expiration time to determine whether said

10 lock is currently granted.

28. The computer program product of Claim 27, further comprising:

executable code that sends, by said current lock owner to said requesting node, a lock response message indicating whether said lock request message has been granted or

15 denied.

29. The computer program product of Claim 28, further comprising:

executable code that determines by said requesting node, that said lock request message has been denied if one of: no lock response message is received from said lock owner in accordance with a timeout condition, and a lock response message is received indicating that said lock request message has been denied.

20

30. The computer program product of Claim 29, wherein said requesting node is granted said lock until a lock expiration time is reached, said requesting node becomes unavailable, and said lock expires automatically when said lock expiration time is reached.

5

31. The computer program product of Claim 29, wherein said current lock owner becomes unavailable as a lock owner, and the computer program product further comprises executable code that performs, by a second node:

determining a new current lock owner of the lock in accordance with an indicator specifying which one or more nodes are available as lock owners, said indicator indicating that said current lock owner is not available; and

sending a lock request message to said new current lock owner, wherein said determining provides for dynamic reassignment of a new lock owner of said lock.

15 32. The computer program product of Claim 31, wherein said distributed lock management is for a distributed lock management in a data storage system including a plurality of nodes, each of said plurality of nodes being a director in said data storage system.

20 33. The computer program product of Claim 32, wherein said lock is associated with a resource used by two or more nodes in said data storage system.

34. The computer program product of Claim 33, wherein said resource is a
portion of global memory.

35. The computer program product of Claim 33, wherein each of said lock
5 request message and said lock response message are sent using a message switch
included in said data storage system.

36. The computer program product of Claim 35, wherein said lock request
message and said lock response message are transmitted between nodes in said data
10 storage system using portions of a message switch used in transmitting data messages.

37. The computer program product of Claim 32, wherein each of said plurality of
nodes indicated as being available as a lock owner in accordance with said indicator is
designated as a lock owner of a lock.
15

38. The computer program product of Claim 37, wherein any one of said plurality
of nodes may request any lock, and wherein any one of said plurality of nodes may be
designated as a lock owner of any lock in accordance with said indicator.

39. The computer program product of Claim 23, wherein said distributed lock management is used in a data storage system including a plurality of nodes, said plurality of nodes being directors that are endpoint in said data storage system, and said executable
5 code that determines and said executable code that sends are included in each of said plurality of nodes and executed prior to sending a lock request message.

40. A computer program product for distributed lock management in a data storage system comprising:

executable code that determines, by a requesting node requesting a lock, a current lock owner of the lock in accordance with an indicator specifying which one or more

5 nodes are available as lock owners; and

executable code that sends a lock request message to said current lock owner, and wherein said executable code that determines is executed by said requesting node prior to executing said executable code that sends for each request.

10 41. The computer program product of Claim 40, wherein said data storage system includes a plurality of nodes which are processors, each of said plurality of nodes being an endpoint and being a lock owner as designated by said indicator when said each node is available as a lock holder.

15 42. The computer program product of Claim 40, wherein said lock request message includes a time period for which said requesting node is requesting said lock, and wherein said current lock owner determines an automatic expiration period for said lock using said time period.

20

43. The computer program product of Claim 42, wherein said current lock owner includes executable code that determines whether to grant or deny said lock request message in accordance with a current value for said automatic expiration time period determined in accordance with a prior lock request message.

5